

Stackelberg Security Games (SSG) Basics and Application Overview

Bo An and Milind Tambe

21.1 Introduction

Security is a critical concern around the world that arises in protecting our ports, airports, transportation or other critical national infrastructure, curtailing the illegal flow of drugs, weapons, and money, and suppressing urban crime, as well as in protecting wildlife, fish, and forests from poachers and smugglers; and it arises in problems ranging from physical to cyber-physical systems. In all of these challenges, we have limited security resources that prevent full security coverage at all times. Instead, limited security resources must be deployed intelligently, taking into account differences in priorities of targets requiring security coverage, the responses of adversaries to the security posture, and potential uncertainty over the types, capabilities, knowledge, and priorities of adversaries faced.

Game theory is well suited to adversarial reasoning for security resource allocation and scheduling problems. Casting the problem as a Bayesian Stackelberg game in consideration of uncertainties, we have developed new algorithms for efficiently solving such games to provide randomized patrolling or inspection strategies. These algorithms have led to some initial successes in this arena, leading to advances over previous approaches in security resource scheduling and allocation, e.g., by addressing key weaknesses of predictability of human schedulers. These algorithms are now deployed in multiple applications: ARMOR has been deployed at the Los Angeles International Airport (LAX) since 2007 to randomize checkpoints on the roadways entering the airport and canine patrol routes within the airport terminals (Pita et al., 2008); IRIS is a game-theoretic scheduler for randomized deployment of the U.S. Federal Air Marshal Service (FAMS) requiring significant scale-up in underlying algorithms that has been in use since 2009 (Tsai, Rathi, Kiekintveld, Ordóñez, & Tambe, 2009); and PROTECT, which requires further scale-up, is deployed for generating randomized patrol schedules for the U.S. Coast Guard in Boston, New York, Los Angeles, and other ports around the United States (An et al., 2013; An, Pita,

Shieh, Tambe, Kiekintveld, & Marecki, 2011; Shieh et al., 2012). Furthermore, TRUSTS is being evaluated for deterring fare evasion, suppressing urban crime, and counterterrorism within the Los Angeles Metro System (Jiang, Nguyen, Tambe, & Procaccia, 2013; Zin & Tambe, 2012), and GUARDS was earlier tested by the U.S. Transportation Security Administration (TSA) for security inside the airport (Pita, Tambe, Kiekintveld, Cullen, & Steigerwald, 2011). Moreover, recent work on “green security games” has led to testing our decision aids for protection of fisheries with the U.S. Coast Guard and protection of wildlife at sites in multiple countries, and “opportunistic crime security games” have focused on suppressing urban crime. These initial successes point the way to major future applications in a wide range of security arenas, with major research challenges in scaling up our game-theoretic algorithms, to addressing human adversaries’ bounded rationality and uncertainties in action execution and observation, as well as in preference elicitation and multi-agent learning.

This chapter will discuss applications of security games, and outline research challenges in security games, including algorithms for scaling up security games, as well as for handling significant adversarial uncertainty and learning models of human adversary behaviors. The rest of this chapter is organized as follows. We start with introducing Stackelberg security games (SSG) basics in Section 21.2. Section 21.3 categorizes different types of security game models. Section 21.4 discusses both deployed and emerging security applications. Section 21.5 overviews some applications of security games beyond the security domains. Section 21.6 concludes with open research issues and future research directions.

21.2 Stackelberg Security Games (SSG) Basics

A generic Stackelberg game has two players, a *leader* and a *follower*. A leader commits to a strategy first, and then a follower optimizes her reward, considering the action chosen by the leader (von Stengel & Zamir, 2004). The two players in a Stackelberg game need not represent individuals, but could also be groups that cooperate to execute a joint strategy, such as a police force or a terrorist organization. Each player has a set of possible *pure strategies*, or the actions that she can execute. A *mixed strategy* allows a player to play a probability distribution over pure strategies. Payoffs for each player are defined over all possible pure-strategy outcomes for both the players. The payoff functions are extended to mixed strategies by taking the expectation over pure-strategy outcomes. The follower can observe the leader’s strategy, and then act in a way to optimize his own payoffs.

To see the advantage of being the leader in a Stackelberg game, consider the game with the payoffs as shown in Table 21.1. The leader is the row player

Table 21.1. *Payoff table for example Stackelberg game*

	c	d
a	3,1	5,0
b	2,0	4,2

and the follower is the column player. The only pure-strategy Nash equilibrium for this game is when the leader plays a and the follower plays c, which gives the leader a payoff of 3; in fact, for the leader, playing b is strictly dominated.

However, in this game, if the leader can commit to playing b before the follower chooses his strategy, then the leader will obtain a payoff of 4, since the follower would then play d to ensure a higher payoff for himself. If the leader commits to a mixed strategy of playing a and b with equal (0.5) probability, then the follower will play d, leading to a higher expected payoff for the leader of 4.5. As we can see from this example, the equilibrium strategy in the Stackelberg game can be in fact different from the Nash equilibria.

Stackelberg games are used to model the attacker-defender strategic interaction in security domains, and this class of Stackelberg games (with certain restrictions on payoffs) (Yin, Korzhyk, Kiekintveld, Conitzer, & Tambe, 2010) is called *Stackelberg security games*. In the Stackelberg security game framework, the security force (defender) is modeled as the leader and the terrorist adversary (attacker) is in the role of the follower. The defender commits to a mixed (randomized) strategy, whereas the attacker conducts surveillance of these mixed strategies and responds with a pure strategy of an attack on a target. Thus, the Stackelberg game framework is a natural approximation of real-world security scenarios. In contrast, the surveillance activity of the attacker cannot be modeled in the simultaneous move games with the Nash equilibrium solution concept. The objective is to find the optimal mixed strategy for the defender.

21.3 Categorizing Security Games

With progress in security games research and the expanding set of applications, it is valuable to consider categorizing this work into three separate areas. These categories are driven by applications, but they also impact the types of games (e.g., single-shot vs. repeated games) considered and the research issues that arise. Specifically, the three categories are:

- (i) infrastructure security games; (ii) green security games; (iii) opportunistic crime security games. We discuss each category next.

21.3.1 Infrastructure Security Games

This type of games and their applications are where the original research on security games was initiated. Key characteristics of these games include the following:

- *Application characteristics:* These games are focused on applications of protecting infrastructure, such as ports, airports, trains, flights, and so on; the goal is often assisting agencies engaged in counterterrorism. Notice that the infrastructure being protected tends to be static, and little changes in a few months, e.g., an airport being protected may have new construction once in two or three years. The activities in the infrastructure are regulated by well-established schedules of movement of people or goods. Furthermore, the targets being protected often have a discrete structure, e.g., terminals at an airport, individual flights, individual trains, etc.
- *Overall characteristics of the defender and adversary play:* These games are single-shot games. The defender does play her strategy repeatedly, i.e., the defender commits to a mixed strategy in this security game. This mixed strategy may get played for months at a time. However, a single attack by an adversary ends the game. The game could potentially restart after such an attack, but it is not set up as a repeated game as in the game categories described later in this chapter.
- *Adversary characteristics:* The games assume that the adversaries are highly strategic and that they may attack after careful planning and surveillance. These carefully planned attacks have high consequences. Furthermore, since these attacks are a result of careful planning with the anticipation of high consequences, attackers commit to these plans of attacks and are not considered to opportunistically move from target to target.
- *Defender characteristics:* The defender does not repeatedly update her strategies. In these domains, there may be just a few attacks that may occur, but these tend to be rare; there are not a very large number of attacks that occur repeatedly. As a result, traditionally, no machine learning is used for the defender to update her strategies over time.

21.3.2 Green Security Games

This type of games and their applications are focused on trying to protect the environment; we adopt the term from “green criminology.”¹

- *Application characteristics:* These games are focused on applications of protecting the environment, including forests, fish, and wildlife. The goal is thus often to assist security agencies against poachers, illegal fishermen, or those

illegally cutting trees in national parks in countries around the world. Unlike infrastructure security games, animals or fish being protected may move around in geographical space, introducing new dimensions of complexity. Finally, the targets being protected are spread out over vast, open geographical spaces, e.g., protecting trees from illegal cutting in large forest regions.

- *Overall characteristics of the defender and adversary play:* These games are not single-shot games. Unfortunately, the adversary often conducts multiple repeated “attacks,” e.g., poaching animals repeatedly. Thus, a single illegal activity does not end the game. Instead, after obtaining reports, e.g., over a month of illegal activities, the defender often re-plans her security activities. In other words, these are repeated security games where the defender plays a mixed strategy, while the attacker attacks multiple times, and then the defender re-plans and plays a new mixed strategy and the cycle repeats. Notice also that the illegal activities of concern here may be conducted by multiple individuals, and thus multiple adversaries are active at any one point.
- *Adversary characteristics:* As mentioned earlier, the adversaries are engaged in repeated illegal activities; the consequences of failure or success are not as severe as in the case of counterterrorism. As a result, every single attack (illegal action) cannot be carried out with the most detailed surveillance and planning. The adversaries will hence exhibit more bounded rationality and bounded surveillance in these domains. Nonetheless, these domains are not ones where illegal activities can be conducted opportunistically (as in the opportunistic crime security games discussed later). This is because in these green security games, the adversaries often have to act in extremely dangerous places (e.g., deep in forests, protecting themselves from wild animals), and thus given the risks involved, they cannot take an entirely opportunistic approach.
- *Defender characteristics:* Since this is a repeated game setting, the defender repeatedly updates her strategies. Machine learning can now be used in this work for the defender to update her strategies over time, given that attack data are available over time. The presence of large amounts of such attack data is very unfortunate in that very large numbers of crimes against the environment are recorded in real life, but the silver lining is that the defender can improve her strategy exploiting these data.

21.3.3 Opportunistic Crime Security Games

This type of games and their applications are focused on trying to combat opportunistic crime. Such opportunistic crime may include criminals engaged

in thefts such as snatching cell phones in metros or stealing student laptops from libraries.

- *Application characteristics*: These games focus on applications involving protecting the public against opportunistic crime. The goal is thus often to assist security agencies in protecting public's property such as cell phones, laptops, or other valuables. Here, human crowds may move around based on scheduled activities, e.g., office hours in downtown settings, or class timings on a university campus, and thus the focus of what needs to be protected may shift on a regular schedule. At least in urban settings, these games focus on specific, limited geographical areas as opposed to vast, open spaces as involved in "green security games."
- *Overall characteristics of the defender and adversary play*: While these games are not explicitly formulated as repeated games, the adversary may conduct or attempt to conduct multiple "attacks" (thefts) in any one round of the game. Thus, the defender commits to a mixed strategy, but a single attack by a single attacker does not end the game. Instead, multiple attackers may be active at a time, conducting multiple thefts while the defender attempts to stop these thefts from taking place.
- *Adversary characteristics*: Once again, the adversaries are engaged in repeated illegal activities, and the consequences of failure or success are not as severe as in the case of counterterrorism. As a result, given that every single attack (illegal action) cannot be carried out with the most detailed surveillance and planning, the adversaries may act even less strategically, and exhibit more bounded rationality and bounded surveillance in these domains. Furthermore, the adversaries are not as committed to detailed plans and are flexible in their execution of their plans, as targets of opportunity present themselves.
- *Defender characteristics*: Available crime data can be used to aid the defender in planning and adapting the defense strategy in response to the criminals' actions. Machine-learning techniques are particularly applicable here.

Even though we have categorized the research and applications of security games in these three categories, not everything is very clearly divided in this fashion. Further research may reveal other categories or generate subcategories of these three categories.

21.4 Deployed and Emerging Security Applications

As we discussed at the beginning of this chapter, the past several years have witnessed the successful application of multi-agent systems in allocating

Table 21.2. *Deployed key applications and associated research problems*

Major Applications	Research challenges and major references
ARMOR for Los Angeles International Airport	Algorithms for scheduling check points and canine units (Paruchuri, Pearce, Marecki, Tambe, Ordóñez, & Kraus, 2008; Pita et al., 2008)
IRIS for U.S. Federal Air Marshals Service	Algorithms for solving large-scale games with complex scheduling constraints (Jain, Kardes, Kiekintveld, Ordóñez, & Tambe, 2010; Tsai et al., 2009)
PROTECT for U.S. Coast Guard	Scheduling boat patrol, modeling human behavior (An et al., 2013; Fang, Jiang, & Tambe, 2013; Shieh et al., 2012)
GUARDS for U.S. Transportation Security Agency	Dealing with heterogeneous defender activities (Pita et al, 2011)
TRUSTS for Urban Security in Transit Systems	Deal with spatial and temporal travel constraints and interruptions (Jiang et al., 2013; Zin & Tambe, 2012)
PAWS for protecting wildlife	Learning behavior model from poaching data (Fang, Stone, & Tambe, 2015; Yang, Ford, Tambe, & Lemieux, 2014)
Protecting public events	Continuous and infinite strategy space (Yin, An, & Jain, 2014; Yin, Xu, Gan, An, & Jiang, 2015)

limited resources to protect critical infrastructures (An, Brown, Vorobeychik, & Tambe, 2013; An, Jain, Tambe, & Kiekintveld, 2011; An et al., 2013; An et al., 2011; Basilico, Gatti, & Amigoni, 2009; Jain, Tsai, Pita, Kiekintveld, Rathi, Tambe, & Ordóñez, 2010; Jiang et al., 2013; Korzhyk, Conitzer, & Parr, 2010; Pita et al., 2011; Tambe & An, 2012; Zin & Tambe, 2012). In the rest of this section, we describe in detail the application of the Stackelberg game framework in multiple significant security domains. Table 21.2 outlines different applications and associated research problems, as well as major references.

21.4.1 Infrastructure Security Game Applications

ARMOR for Los Angeles International Airport

Los Angeles International Airport (LAX) is the largest destination airport in the United States and serves 60–70 million passengers per year. The LAX police use diverse measures to protect the airport, which include vehicular check-points and police units patrolling with canines. The eight different terminals at LAX have very different characteristics, like physical size, passenger loads, foot traffic, or international versus domestic flights. Furthermore, the numbers

of available vehicle checkpoints and canine units are limited by resource constraints. Thus it is challenging to optimally allocate these resources to improve their effectiveness while avoiding patterns in the scheduled deployments.

The Assistant for Randomized Monitoring over Routes (ARMOR) system focuses on two of the security measures at LAX (checkpoints and canine patrols) and optimizes security resource allocation using Bayesian Stackelberg games. Take the vehicle checkpoints model as an example. Assume that there are n roads; the police's strategy is placing $m < n$ checkpoints on these roads where m is the maximum number of checkpoints. The adversary may choose to attack through one of these roads. ARMOR models different types of attackers with different payoff functions, representing different capabilities and preferences for the attacker. ARMOR uses Decomposed Optimal Bayesian Stackelberg Solver (DOBSS) to compute the defender's optimal strategy (Paruchuri et al., 2008). ARMOR has been successfully deployed since August 2007 at LAX to randomize checkpoints on the roadways entering the airport and canine patrol routes within the airport terminals (Pita et al., 2008).

IRIS for U.S. Federal Air Marshals Service

The U.S. Federal Air Marshals Service (FAMS) allocates air marshals to flights originating in and departing from the United States to dissuade potential aggressors and prevent an attack. Flights are of different importance based on a variety of factors such as the numbers of passengers, the population of source/destination, international flights from different countries, and special events that can change the risks for particular flights at certain times. Security resource allocation in this domain is significantly more challenging than for ARMOR: a limited number of FAMS need to be scheduled to cover thousands of commercial flights each day. Furthermore, these FAMS must be scheduled on tours of flights that obey various constraints (e.g., the time required to board, fly, and disembark). Therefore, we face significant computational challenge while generating the optimal scheduling policy that meets these scheduling constraints.

Against this background, the Intelligent Randomization in Scheduling (IRIS) system has been developed and has been deployed by FAMS since October 2009 to randomize schedules of air marshals on international flights. In IRIS, the targets are the set of n flights and the attacker could choose to attack one of these flights. FAMS can assign $m < n$ air marshals to protect these flights. Since the number of possible schedules exponentially increases with the number of flights and resources, DOBSS is no longer applicable to the FAMS domain. Instead, IRIS uses the much faster ASPEN algorithm (Jain et al., 2010) to generate the schedule for thousands of commercial flights per day.



Figure 21.1. USCG boats patrolling the ports of Boston and NY. (a) PROTECT is being used in Boston. (b) Extend PROTECT to NY.

IRIS also uses an attribute-based preference elicitation system to determine reward values for the Stackelberg game model.

PROTECT for U.S. Coast Guard

The U.S. Coast Guard's (USCG) mission includes maritime security of the U.S. coasts, ports, and inland waterways – a security domain that faces increased risks due to threats such as terrorism and drug trafficking. Given a particular port and the variety of critical infrastructure that an adversary may attack within that port, the USCG conducts patrols to protect this infrastructure; however, while the adversary has the opportunity to observe patrol patterns, limited security resources imply that USCG patrols cannot be at every location 24/7. To assist the USCG in allocating its patrolling resources, the Port Resilience Operational/Tactical Enforcement to Combat Terrorism (PROTECT) model is being designed to enhance maritime security and has been in use at the port of Boston since April 2011 (Figure 21.1). Similar to previous applications ARMOR and IRIS, PROTECT uses an attacker-defender Stackelberg game framework, with the USCG as the defender against terrorist adversaries that conduct surveillance before potentially launching an attack.

PROTECT is currently deployed in the ports of Boston, New York, Los Angeles/Long Beach, and several others (An et al., 2013). Indeed, the goal now is to deploy PROTECT at ports nationwide. Furthermore, beyond just port protection, PROTECT has been extended to protect ferry systems (Figure 21.2) such as the Staten Island ferry in New York (Fang et al., 2013).

While PROTECT builds on previous work, it offers some key innovations. First, to improve PROTECT's efficiency, a compact representation of the defender's strategy space is used by exploiting equivalence and dominance.



Figure 21.2. Protecting ferries with patrol boats.

Second, the evaluation of PROTECT for the first time provides real-world data: (i) comparison of human-generated vs PROTECT security schedules; and (ii) results from an adversarial perspective team's (human mock attackers) analysis.

GUARDS for U.S. Transportation Security Agency

The U.S. Transportation Security Administration (TSA) is tasked with protecting the nation's more than 400 airports. To aid the TSA in scheduling resources to protect airports, a new application called Game-theoretic Unpredictable and Randomly Deployed Security (GUARDS) has been developed. While GUARDS also utilizes Stackelberg games as ARMOR and IRIS, GUARDS faces three key challenges Pita et al., 2011: 1) reasoning about hundreds of heterogeneous security activities; 2) reasoning over diverse potential threats; and 3) developing a system designed for hundreds of end-users. To address those challenges, GUARDS created a new game-theoretic framework that allows for heterogeneous defender activities and compact modeling of a large number of threats and developed an efficient solution technique based on general-purpose Stackelberg game solvers. GUARDS was originally tested at an undisclosed airport and further results are awaited (Pita et al., 2011).

TRUSTS for Urban Security in Transit Systems

TRUSTS focuses on three major security challenges: deterring fare evasion, suppressing crime, and counterterrorism. Significant focus in TRUSTS has been on



Figure 21.3. TRUSTS for transit systems.

detering fare evasion. Specifically, in some urban transit systems, including the Los Angeles Metro Rail system, passengers are legally required to purchase tickets before entering, but are not physically forced to do so (Figure 21.3). Instead, patrol units move about through the transit system, inspecting tickets of passengers, who face fines for fare evasion. This setting yields the problem of computing optimal patrol strategies to deter fare evasion and hence maximize revenue. The Tactical Randomization for Urban Security in Transit Systems (TRUSTS) system models the patrolling problem as a leader-follower Stackelberg game (Jiang et al., 2013; Zin & Tambe, 2012). Urban transit systems, however, present unique computational challenges since there are exponentially many possible patrol strategies, each subject to both the spatial and temporal constraints of travel within the transit network under consideration. To overcome this challenge, TRUSTS uses a compact representation that captures the spatial as well as temporal structure of the domain. The system has been evaluated using real-world ridership data from the Los Angeles Metro Rail system.

One key finding from initial tests was that the schedules generated by officers were often interrupted. Interruptions occurred because in frequent interactions with the public, sometimes officers would get stopped by lost travelers, or sometimes they would need to arrest someone. Such interruptions mean that the schedules now need to be highly dynamic. To that end, a new generation of Stackelberg game-based scheduling algorithms – using Markov Decision Problems – was designed. This led to schedules now being loaded onto smartphones and given to officers. The schedules are then automatically updated on the smartphones if interruptions occur (Luber, Yin, Fave, Jiang, Tambe, & Sullivan, 2013).

Protecting Public Events

Security games can also be applied to resource allocation for protecting public events. Public events in major cities are prime terrorism targets since they usually provide easy access to a large number of targets for the adversary. There have been some successful terrorist attacks on large, public events in the United States and Europe in the past few years, e.g., the recent Boston Marathon bombings on April 15, 2013, and the July 7, 2005, London bombings. Intelligent deployment of limited security resources to protect such events is therefore extremely important and challenging since the importance of targets changes over time. For example, the value of targets along a marathon track changes over time with the changing number of participants and spectators at any specific area over the course of the race. In addition, since the attacker may attack at any time and the defender can relocate resources among targets at any time, the strategy space of each agent is continuous and infinite. Furthermore, due to the relative infrequency of such events, the attacker does not get the opportunity to conduct surveillance and respond to a distribution of defender strategies. In this case, a pure defender strategy sampled from the optimal mixed strategy does not necessarily outperform the one-shot optimal pure strategy in terms of ex-post payoff. Algorithms have been proposed to compute the optimal pure defender strategy despite the infinite strategy space and time-varying target values (Yin et al., 2014; Yin et al., 2015).

21.4.2 Green Security Game Applications

A number of newer applications are focused on suppressing environmental crime. One of those is protecting forests (Johnson, Fang, Yang, Tambe, & Albers, 2012), where we must protect a continuous forest area from extractors. Since the attacker's behavior (e.g., extracting important resources from the forest) could be effected by spatial considerations, it is critical for the defender to incorporate spatial considerations into her enforcement decisions (Albers, 2010).

Another area of interest is protecting endangered species. Endangered species poaching is reaching critical levels as the populations of these species plummet to unsustainable numbers. The global tiger population, for example, has dropped more than 95% from the start of the 1900s and has resulted in three out of nine species' extinctions. Depending on the area and animals poached, motivations for poaching range from profit to sustenance, with the former being more common when profitable species such as tigers, elephants, and rhinos are the targets. To counter poaching efforts and to rebuild the species' populations, countries have set up protected wildlife reserves and conservation agencies



Figure 21.4. Computational game theory can play a role in suppressing environmental crime.

tasked with defending these large reserves. Because of the size of the reserves and the common lack of law enforcement resources, conservation agencies are at a significant disadvantage when it comes to deterring and capturing poachers. Agencies use patrolling as a primary method of securing the parks. Due to their limited resources, however, patrol managers must carefully create patrols that account for many different variables (e.g., limited patrol units to send out, multiple locations that poachers can attack at varying distances to the outpost). Protection Assistant for Wildlife Security (PAWS) aims to assist conservation agencies in their critical role of patrol creation by predicting where poachers will attack and optimizing patrol routes to cover those areas (Figure 21.4; Yang et al., 2014).

Another emerging application domain is that of ensuring the sustainability of fish resources. Marine fisheries are acknowledged to be some of the most important food resources for countries around the world. As reported by the World Wildlife Fund for Nature (WWF), cod are currently at risk from overfishing in the UK, Canada, and most other Atlantic countries. Global cod catch has suffered a 70% drop over the past 30 years, and, if this trend continues, the world's cod stocks will disappear in 15 years. Illegal, unreported, and unregulated (IUU) fishing is one of the major threats to the sustainability of ocean fish resources. As estimated by the National Oceanic and Atmospheric Administration (NOAA), IUU fishing produces between 11 million and 26 million tons of seafood annually, representing as much as 40% of the total catch in some fisheries. The driver behind IUU fishing is high economic profit and low chance of seizure. It is impossible



Figure 21.5. U.S. Coast Guard personnel on a mission to protect fisheries.

to maintain a 24/7 presence to prevent IUU fishing everywhere due to the limited patrolling resources. Hence the allocation of the patrolling resources becomes a key challenge for security agencies like the U.S. Coast Guard (Figure 21.5).

21.4.3 Opportunistic Crime Security Game Applications

A notable characteristic of urban crime, distinct from organized terrorist attacks, is that most urban crimes are opportunistic in nature, i.e., criminals do not plan their attacks in detail; rather, they seek opportunities for committing crime and are agile in their execution of the crime (Short, D’Orsogna, Pasour, Tita, Brantingham, Bertozzi, & Chayes, 2008; Zhang, Jiang, Short, Brantingham, & Tambe, 2014). Thus, the Opportunistic Crime Security Game (OCSG) model is a good fit for addressing crime in urban settings.

A particular category of urban crime is crime on transit systems such as phone snatching.

OCSG has been validated in trials on the LA Metro systems with good results (Fave et al., 2014). The trials ran for two days with each test consisting of a two-hour patrol involving two teams of two security officers. Each team had to patrol seven stations of a particular LA Metro train line using schedules generated using the OCSG framework. Figure 21.6 shows such probabilities and correlates them to the crime statistics for each of the 14 stations to patrol. In the figure, the x-axis enumerates the 14 stations to patrol. The bar graphs (y-axis on the right) show, for each station, the total number of crimes that happened during 2012 and 2013. Finally, the line graph shows the different coverage probabilities calculated for each station (y-axis on the left). In the figure, the stations with a larger coverage probability (stations 5 to 10) are either the stations with a large number of crimes (stations 5 and 8) or the adjacent stations (stations 6, 7, 9, and 10). The latter stations are given a large coverage probability because the OCSG model anticipates the possibility that criminals will choose stations 6, 7, 9, and 10 anticipating that stations 5 and 8 will be frequently patrolled by security officers. Hence, these coverage probabilities show how the OCSG




Figure 21.6. Crime statistics and coverage probabilities.

model allows schedulers to build effective, real-world patrol schedules. During the tests, the officers were able to write five citations and make two arrests. In general, they were able to understand and follow the schedule easily.

There are many other possible applications of OCSG, including scheduling police patrols and providing security at crowded places such as fairs.

21.5 Applications of Security Games beyond Security

The security games model is rich enough to capture many real-world scenarios of defender and adversary interaction beyond the physical security settings discussed here. One interesting work, called audit games (Blocki, Christin, Datta, Procaccia, & Sinha, 2013, 2015) enhances the security games model with choice of punishments in order to capture scenarios of security and privacy policy enforcement in large organizations. Large organizations (such as Google, Facebook, and hospitals) hold enormous amounts of privacy-sensitive data. These organizations mandate their employees to adhere to certain privacy policies when accessing data. Auditing of access logs is used by organizations to check for policy-violating accesses and then the violators are punished. Auditing often requires human help to investigate suspicious cases, and thereby arises the problem of allocating few resources to the huge number of cases to investigate. Another relevant question in this domain is how much should the organization punish in case of a violation?

The audit game models the adversary as an agent that performs certain tasks (e.g., accesses to private data), and a subset of these tasks is policy violations. The auditor inspects a subset of the tasks, and detects violations from the inspected set. As punishments do affect the behavior of the adversary, it is critical for the auditor to choose the right level of punishment. As a consequence, the choice of a punishment level is added to the action space of the auditor. However, punishment is not free for the auditor, the intuition being that a high punishment level creates a hostile work environment, leading to lack in

productivity of employees that results in loss for the organization (auditor). As a consequence, the auditor cannot impose infinite punishment and deter any adversary. The auditor's cost for a punishment level is modeled as a loss proportional to the choice of the punishment level. The auditor moves first by committing to an inspection and punishment strategy, followed by the best response of the adversary. The resultant Stackelberg equilibrium optimization turns out to be non-convex due to the punishment variable. The authors present efficient algorithms for various types of scheduling constraints.

There are various other interesting applications such as using a security game model to randomize regression testing in the domain of software testing (Kukreja, Halfond, & Tambe, 2013). In software testing, due to time constraints, often not all the tests for a software product can be run. Thus, this leads to the problem of using few resources (test) to protect against attacks (software bugs); the security game model is a natural fit for such a scenario. Another interesting application (Li & Conitzer, 2013) is in the domain of testing students. Often the set of all exam questions for large-scale exams (e.g., driver's license test) is known a priori. However, if this set is large, the test takers cannot memorize all the answers. Then, the examiner must choose questions randomly to best test the students on the questions-answers they did not memorize. This can also be posed as a Stackelberg game. Other applications include the scenario of scheduling randomized patrols for improving security of transit networks in Singapore (Varakantham, Lau, & Yuan, 2013).

21.6 Major Research Issues

In this section, we highlight some key research challenges, including scalability, robustness, and human adversary modeling. Readers can refer to other chapters (e.g., basic solution concepts and algorithms, linear programming methods, addressing the unpredictable human element, and learning) of this book for detailed discussions of the research issues outlined next.

Scalability: The first research challenge is improving the scalability of our algorithms for solving Stackelberg (security) games. The strategy space of both the defender and the attacker in these games may exponentially increase with the number of security activities, attacks, and resources. As we scale up to larger domains, it is critical to develop newer algorithms that scale up significantly beyond the limits of the current state of the art of Bayesian Stackelberg solvers. Driven by the growing complexity of applications, a sequence of algorithms for solving security games has been developed, including DOBSS (Paruchuri et al., 2008), ERASER (Jain et al., 2010), and ASPEN (Jain et al., 2010). However, existing algorithms still cannot scale up to very large-scale

domains such as scheduling randomized checkpoints in cities. In such graph-based security games, the strategy space of the defender grows exponentially with the number of available resources and the strategy space of the attacker grows exponentially with the size of the road network considered. Three key ideas have been used to design efficient algorithms, including (i) Marginals (compact representation); (ii) Incremental strategy generation; and (iii) Incremental constraint generation. For patrolling domains with spatiotemporal constraints, defender mixed strategies can be compactly represented as fractional flows. This approach has recently been applied to efficiently compute fare-enforcement patrols in urban transit systems (Jiang, Yin, Kraus, Zhang, & Tambe, n.d.; Zin & Tambe, 2012) and boat patrols for protecting ferries (Fang et al., 2013). To schedule checkpoints, an incremental strategy generation approach called “double oracle” is applied that does not require the enumeration of the entire strategy space for either of the players (Jain, Korzhlyk, Vanek, Pechoucek, Conitzer, & Tambe, 2011; Jain, Tambe, & Conitzer, n.d.). When solving large-scale security games, we may need to deal with large (even infinite) set of constraints and the constraint generation approach was recently applied (Nguyen, Yadav, An, Tambe, & Boutilier, 2014). The key idea is to sample a subset of constraints and gradually expand this set by adding violated constraints to the relaxed problem until convergence to the optimal solution.

Robustness: The second challenge is improving solutions’ robustness. Classical game theory solution concepts often make assumptions on the knowledge, rationality, and capability of players. Unfortunately, those assumptions could be wrong in real-world scenarios. Therefore, while computing the defender’s optimal strategy, algorithms should take into account various uncertainties faced in the domain, including payoff noise (Kiekintveld, Marecki, & Tambe, 2011), execution/observation error (Yin, Jain, Tambe, & Ordóñez, 2011), and uncertain capability (An, Tambe, Ordóñez, Shieh, & Kiekintveld, 2011). For observation uncertainty, it is typically assumed that the attacker has perfect knowledge of the defender’s randomized strategy or can learn the defender’s strategy after conducting a fixed period of surveillance. In consideration of surveillance cost, these assumptions are clearly simplistic since attackers may act with partial knowledge of the defender’s strategies and may dynamically decide whether to attack or conduct more surveillance. Security game models with limited observation (An et al., 2013; An, Kempe, Kiekintveld, Shieh, Singh, Tambe, & Vorobeychik, 2012) have been proposed in which the attacker either makes limited number of observations or dynamically determines a place to stop surveillance. Since the belief state space exponentially increases with observation length, it is still computationally challenging to solve large games in consideration of limited observation.

Bounded Rationality: One required research direction with respect to robustness is addressing bounded rationality of human adversaries, which is a fundamental problem that can affect the performance of our game-theoretic solutions. Recently, there has been some research on applying ideas (e.g., prospect theory (Kahneman & Tvesky, 1979) and quantal response (McKelvey & Palfrey, 1995)) from social science or behavioral game theory within security game algorithms (Nguyen, Yang, Azaria, Kraus, & Tambe, 2013; Pita, Jain, Tambe, Ordóñez, & Kraus, 2010; Yang, Jiang, Tambe, & Ordóñez, 2013; Yang, Kiekintveld, Ordóñez, Tambe, & John, 2011). Previous work usually applies existing frameworks and sets the parameters of these frameworks by experimental tuning or learning. However, in real-world security domains, we may have very limited data, or may have only limited information on the biases displayed by adversaries. Recently, monotonic maximin (Jiang et al., 2013) was proposed as a robust solution concept to Stackelberg security games with boundedly rational adversaries. It tries to optimize defender utility against the worst-case monotonic adversary behavior, where monotonicity is the property that actions with higher expected utility are played with higher probability. An open research challenge is to combine such robust-optimization approaches with available behavior data. Furthermore, since real-world human adversaries are sometimes distributed across coalitions of socially, culturally, and cognitively biased agents acting behind a veil of uncertainty, we may need significant interdisciplinary research to build in social, cultural, and cognitive biases into our adversary models.

Planning and Learning: The challenges of planning and learning pertain to new security domains such as patrolling metro systems and forests. In these domains, schedules might be interrupted due to some unexpected events (e.g., writing a citation in a train line or boarding an illegal fisherman's boat). As a consequence, incorporating planning within security games is necessary to model stochastic decision making. Furthermore, in the real world, a schedule's spatial and temporal constraints are typically continuous dimensions. Hence, more expressive models need to be derived to represent such schedules. So far, data available in most domains have been principally used to define the game's matrices. However, in the newer domains, the planned schedules are frequently interrupted. Hence, the information about the locations and times of interruptions, as well as the information about the actual interactions between attackers and defenders, can also be used. The former can be used to represent the uncertainty of the environment, whereas the latter can be used to learn the attacker's behavior. An online planning algorithm for the protector in resource conservation games is proposed to use the information gained by observing the extractor's actions (Qian, Haskell, Jiang, & Tambe, 2014). The resource conservation game is modeled as a repeated game and then is casted as a POMDP.

Integrating Dynamic Information: The Transportation Security Administration (TSA) is launching a new concept to address its current challenges by integrating and quantifying risk information to comprehensively analyze risk on a per-flight basis. Rather than using a one-size-fits-all approach, the idea is to dynamically tailor the security by integrating information across five threat vectors: passengers, checked baggage, cargo, aircraft operator, and airport/perimeter. This approach will provide a more comprehensive picture of total flight risk, enabling the TSA to reduce risk by allowing for more dynamic allocation of resources. This is a massive project that would radically impact domestic aviation security in the United States. The security games framework is being used as part of this project to quantitatively test and validate the concepts used in this new aviation security concept of the future. This may have a significant impact on how this new concept of dynamic security gets operationalized.

NOTE

1. We use the term *green security games* also to avoid any confusion that may come about given that terms related to the environment and security have been adopted for other uses. For example, the term *environmental security*, broadly speaking, refers to threats posed to humans due to environmental issues, e.g., climate change or shortage of food. The term *environmental criminology*, on the other hand, refers to analysis and understanding of how different environments affect crime.

REFERENCES

- Albers, H. (2010). Spatial modeling of extraction and enforcement in developing country protected areas. *Resource and Energy Economics*, 32(2), 165–179.
- An, B., Brown, M., Vorobeychik, Y., & Tambe, M. (2013). Security games with surveillance cost and optimal timing of attack execution. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (pp. 223–230).
- An, B., Jain, M., Tambe, M., & Kiekintveld, C. (2011). Mixed-initiative optimization in security games: A preliminary report. In *Proceedings of the AAAI Spring Symposium on Help Me Help You: Bridging the Gaps in Human-Agent Collaboration* (pp. 8–11).
- An, B., Kempe, D., Kiekintveld, C., Shieh, E., Singh, S., Tambe, M., & Vorobeychik, Y. (2012). Security games with limited surveillance. In *Proceedings of the 26th Conference on Artificial Intelligence (AAAI)* (pp. 1241–1248).
- An, B., Ordóñez, F., Tambe, M., Shieh, E., Yang, R., Baldwin, C., DiRenzo, J., Moretti, K., Maule, B., & Meyer, G. (2013). A deployed quantal response-based patrol planning system for the U.S. Coast Guard. *Interfaces*, 43(5), 400–420.
- An, B., Pita, J., Shieh, E., Tambe, M., Kiekintveld, C., & Marecki, J. (2011, March). GUARDS and PROTECT: Next generation applications of security games. *SIGECOM*, 10, 31–34.
- An, B., Tambe, M., Ordóñez, F., Shieh, E., & Kiekintveld, C. (2011) Refinement of strong Stackelberg equilibria in security games. In *Proceedings of the 25th Conference on Artificial Intelligence* (pp. 587–593).

- Basilico, N., Gatti, N., & Amigoni, F. (2009). Leader-follower strategies for robotic patrolling in environments with arbitrary topologies. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (pp. 500–503).
- Blocki, J., Christin, N., Datta, A., Procaccia, A. D., & Sinha, A. (2013). Audit games. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence* (pp. 41–47).
- Blocki, J., Christin, N., Datta, A., Procaccia, A. D., & Sinha, A. (2015). Audit games with multiple defender resources. In *AAAI Conference on Artificial Intelligence (AAAI)* (pp. 791–797).
- Fang, F., Jiang, A. X., & Tambe, M. (2013). Optimal patrol strategy for protecting moving targets with multiple mobile resources. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (pp. 957–964).
- Fang, F., Nguyen, T. H., Pickles, R., Lam, W. Y., Clements, G. R., An, B., Singh, A., Tambe, M., & Lemieux, A. (2016). Deploying PAWS: Field optimization of the protection assistant for wildlife security. In *AAAI* (pp. 3966–3973).
- Fang, F., Stone, P., & Tambe, M. (2015). When security games go green: Designing defender strategies to prevent poaching and illegal fishing. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)* (pp. 2589–2595).
- Fave, F. M., Brown, M., Zhang, C., Shieh, E., Jiang, A., Rosoff, H., Tambe, M., & Sullivan, J. (2014). Security games in the field: Deployments on a transit system. In F. Dalpiaz, J. Dix, & M. van Riemsdijk (Eds.), *Engineering multi-agent systems*, volume 8758 of *Lecture Notes in Computer Science*, pages 103–126.
- Jain, M., Kardes, E., Kiekintveld, C., Ordóñez, F., & Tambe, M. (2010). Security games with arbitrary schedules: A branch and price approach. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence* (pp. 792–797).
- Jain, M., Korzhuk, D., Vanek, O., Pechoucek, M., Conitzer, V., & Tambe, M. (2011). A double oracle algorithm for zero-sum security games on graphs. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (pp. 327–334).
- Jain, M., Tambe, M., & Conitzer, V. (n.d.). Security scheduling for real-world networks. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (pp. 215–222).
- Jain, M., Tsai, J., Pita, J., Kiekintveld, C., Rath, S., Tambe, M., & Ordóñez, F. (2010). Software assistants for randomized patrol planning for the LAX airport police and the federal air marshal service. *Interfaces*, 40(4), 267–290.
- Jiang, A. X., Nguyen, T. H., Tambe, M., & Procaccia, A. D. (2013). Monotonic maximin: A robust Stackelberg solution against boundedly rational followers. In *Proceedings of the Conference on Decision and Game Theory for Security (GameSec)* (pp. 119–139).
- Jiang, A., Yin, Z., Kraus, S., Zhang, C., & Tambe, M. (n.d.). Game-theoretic randomization for security patrolling with dynamic execution uncertainty. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (pp. 207–214).

- Johnson, M., Fang, F., Yang, R., Tambe, M., & Albers, H. (2012). Patrolling to maximize pristine forest area. In *Proceedings of the AAAI Spring Symposium on Game Theory for Security, Sustainability and Health*.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291.
- Kiekintveld, C., Marecki, J., & Tambe, M. (2011). Approximation methods for infinite Bayesian Stackelberg games: Modeling distributional uncertainty. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (pp. 1005–1012).
- Korzhyk, D., Conitzer, V., & Parr, R. (2010). Complexity of computing optimal Stackelberg strategies in security resource allocation games. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence* (pp. 805–810).
- Kukreja, N., Halfond, W. G., & Tambe, M. (2013). Randomizing regression tests using game theory. In *Automated Software Engineering (ASE), 2013 IEEE/ACM 28th International Conference on Artificial Intelligence* (pp. 616–621).
- Li, Y., & Conitzer, V. (2013). Game-theoretic question selection for tests. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence* (pp. 254–262).
- Luber, S., Yin, Z., Fave, F. D., Jiang, A. X., Tambe, M., & Sullivan, J. P. (2013). Game-theoretic patrol strategies for transit systems: The trusts system and its mobile app (demonstration). In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)[Demonstrations Track]* (pp. 1377–1378).
- McKelvey, R. D., & Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1), 6–38.
- Nguyen, T., Yadav, A., An, B., Tambe, M., & Boutilier, C. (2014). Regret-based optimization and preference elicitation for Stackelberg security games with uncertainty. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence* (pp. 756–762).
- Nguyen, T. H., Yang, R., Azaria, A., Kraus, S., & Tambe, M. (2013). Analyzing the effectiveness of adversary modeling in security games. In *Conference on Artificial Intelligence (AAAI)* (pp. 718–724).
- Paruchuri, P., Pearce, J. P., Marecki, J., Tambe, M., Ordóñez, F., & Kraus, S. (2008). Playing games with security: An efficient exact algorithm for Bayesian Stackelberg games. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (pp. 895–902).
- Pita, J., Jain, M., Tambe, M., Ordóñez, F., & Kraus, S. (2010). Robust solutions to Stackelberg games: Addressing bounded rationality and limited observations in human cognition. *Artificial Intelligence*, 174(15) 1142–1171.
- Pita, J., Jain, M., Western, C., Portway, C., Tambe, M., Ordóñez, F., Kraus, S., & Parachuri, P. (2008). Deployed ARMOR protection: The application of a game-theoretic model for security at the Los Angeles International Airport. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (pp. 125–132).
- Pita, J., Tambe, M., Kiekintveld, C., Cullen, S., & Steigerwald, E. (2011). GUARDS – game theoretic security allocation on a national scale. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (pp. 37–44).

- Qian, Y., Haskell, W. B., Jiang, A. X., & Tambe, M. (2014). Online planning for optimal protector strategies in resource conservation games. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multiagent Systems* (pp. 733–740).
- Shieh, E., An, B., Yang, R., Tambe, M., Baldwin, C., DiRenzo, J., Maule, B., & Meyer, G. (2012). PROTECT: An application of computational game theory for the security of the ports of the United States. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)* (pp. 2173–2179).
- Short, M. B., D’Orsogna, M. R., Pasour, V. B., Tita, G. E., Brantingham, P. J., Bertozzi, A. L., & Chayes, L. B. (2008). A statistical model of criminal behavior. *Mathematical Models and Methods in Applied Sciences*, 18, 1249–1267.
- Tambe, M., & An, B. (2012). Game theory for security: A real-world challenge problem for multiagent systems and beyond. In *Proceedings of the AAAI Spring Symposium on Game Theory for Security, Sustainability and Health* (pp. 69–74).
- Tsai, J., Rathi, S., Kiekintveld, C., Ordóñez, F., & Tambe, M. (2009). IRIS: A tool for strategic security allocation in transportation networks. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (pp. 37–44).
- Varakantham, P., Lau, H. C., & Yuan, Z. (2013). Scalable randomized patrolling for securing rapid transit networks. In *Conference on Innovative Applications of Artificial Intelligence (IAAI)* (pp. 1563–1568).
- von Stengel, B., & Zamir, S. (2004). Leadership with commitment to mixed strategies. Technical Report LSE-CDAM-2004-01, Centre for Discrete and Applicable Mathematics, London School of Economics and Political Science, London, UK.
- Yang, R., Ford, B. J., Tambe, M., & Lemieux, A. (2014). Adaptive resource allocation for wildlife protection against illegal poachers. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (pp. 453–460).
- Yang, R., Jiang, A. X., Tambe, M., & Ordóñez, F. (2013). Scaling-up security games with boundedly rational adversaries: A cutting-plane approach. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)* (pp. 404–410).
- Yang, R., Kiekintveld, C., Ordóñez, F., Tambe, M., & John, R. (2011). Improving resource allocation strategy against human adversaries in security games. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)* (pp. 458–464).
- Yin, Y., An, B., & Jain, M. (2014). Game-theoretic resource allocation for protecting large public events. In *Proceedings of the 28th Conference on Artificial Intelligence (AAAI)* (pp. 826–834).
- Yin, Y., Jain, M., Tambe, M., & Ordóñez, F. (2011). Risk-averse strategies for security games with execution and observational uncertainty. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence (AAAI)* (pp. 758–763).
- Yin, Y., Xu, H., Gan, J., An, B., & Jiang, A. (2015). Computing optimal mixed strategies for security games with dynamic payoffs. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)* (pp. 681–688).

- Yin, Z., Korzhyk, D., Kiekintveld, C., Conitzer, V., & Tambe, M (2010). Stackelberg vs. Nash in security games: interchangeability, equivalence, and uniqueness. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (pp. 1139–1146).
- Yin, Z., & Tambe, M. (2012). A unified method for handling discrete and continuous uncertainty in Bayesian Stackelberg games. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (pp. 855–862).
- Zhang, C., Jiang, A. X., Short, M. B., Brantingham, P. J., & Tambe, M. (2014). Defending against opportunistic criminals: new game-theoretic frameworks and algorithms. In *Decision and game theory for security* (pp. 3–22). Springer.